Design Considerations For Large Mobile Packet Radio Networks*

Simon S. Lam
Associate Professor of Computer Sciences
University of Texas at Austin
U.S.A.

We consider packet radio networks that are large, have a changing connectivity and support mobile users (key attributes). The primary performance measures are throughput efficiency, network delay and availability/survivability. Four existing routing algorithms are described and compared qualitatively. We concluded that with the above network attributes, these algorithms can provide either throughput efficiency or good availability but none can provide both. A hierarchical network structure and routing strategy are proposed which can conceptually provide both by adapting to the nature of the network. Network transit delays may also be improved with a hierarchical structure. The implementation of such a network architecture requires several classes of packet radios with different transmission power ranges. We also addressed the design of channel access protocols, and the impact of this network architecture on the design of network algorithms for flow control and error control.

* This work was supported in part by U.S. Army CORADCOM, Fort Monmouth, New Jersey and in part by National Science Foundation Grant No. ECS78-01803. The views, opinions and/or findings contained in this report are those of the author and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation. These views, opinions and/or findings also do not reflect the views of the National Science Foundation.

1. INTRODUCTION

The design of packet radio networks with a thousand nodes or more is considered. Some nodes are network users (source/sinks of data packets). Some nodes are network relays, while some user nodes will also serve as relays when needed.

The network connectivity is assumed to change frequently in time due to the presence of mobile users and signal attenuation of radio transmissions for various reasons (fading, line-of-sight obstructions etc.). These changes are said to be passive since they are not under network control. Some changes are subject to network control and are said to be active, such as: (1) transmission power control of relays, and (2) reconfiguration of ground relays. A third attribute that will also impact the design of network algorithms is the diversity of traffic classes and performance requirements that often characterizes very large networks. Data traffic sources may be bursty as well as steady; packet generation statistics of sources may be periodic as well as random. Some messages may require highly reliable transfers while some others are short-lived data and do not need error recovery.

The primary performance requirements to be considered are:
network throughput rate,
network delay, and
availability/survivability.
The network throughput rate is in terms of packets per second transported. The network delay is the time between source input of a packet and its successful delivery to the intended destination; the delay may need to be characterized in terms of its mean, variance or probability distribution. Network availability is a measure of the logical connectivity of the network provided to sources and destinations in the presence of link and relay failures. (Note that physical connectivity does not imply logical connectivity.) The term survivability is used interchangeably to mean network availability when network topology changes are on a large scale. For purposes of this paper, we shall limit our attention to relay and link failures and not worry about the (important) security and anti-jam aspects of networks in a hostile environment.

Additional performance specifications may be given that define detailed network service characteristics. There may be specific throughput requirements and delay constraints for individual source-destination pairs, or for various classes of packets. To satisfy these requirements, network algorithms are needed for fair allocation of resources among demands of equal importance, and priority allocation of resources to important demands. A high level of reliability may sometimes be required for the transport of special classes of messages, necessitating special error control procedures which may impact the network's throughput and delay performance.

In Section 2, we shall examine 4 existing routing algorithms and compare them qualitatively with respect to the above-mentioned performance measures. In Section 3, a hierarchical network structure and routing strategy are proposed. The implementation of such a network architecture requires several classes of packet radios with different transmission power ranges. Its impact on the design of channel access protocols and network algorithms for flow control and error control will also be addressed.

2. NETWORK ROUTING ALGORITHMS

The network routing algorithm plays the central role within the family of network algorithms. We next describe the basic concepts of four routing algorithms that have been considered for packet radio networks [FRAN 75, GITM 76, KAHN 77, KAHN 78]

1. distributed adaptive (ARPANET [MCQU 80]),
2. broadcast routing (flooding),
3. fixed routes established by one or more stations, and
4. fixed routes established by flooding.

The performance characteristics of these routing algorithms are assessed qualitatively with respect to the network attributes and the primary performance requirements discussed earlier. (See Figure 1 on the following page.)

Distributed adaptive routing

A distributed adaptive routing algorithm such as that originally used by ARPANET requires that each relay node in the network maintains a delay table for making routing decisions. Suppose we are considering node i. The delay estimate for node i to destination j by way of neighboring node k is $t_{ik} + d_{kj}$, where t_{ik} is the delay to go from node i to its neighbor node k and can be estimated from the local conditions of node i; d_{kj} is supplied by periodic updates from neighbor k and is node k's minimum delay estimate to go to destination j. A packet residing at node i and destined for node j will be forwarded to the neighbor relay with the smallest delay estimate

$$d_{ij} = \min_k (t_{ik} + d_{kj}).$$

The network throughput efficiency of this routing algorithm is good since approximately shortest delay paths are used for transporting packets. Some overhead is incurred in the periodic exchange of routing table updates between neighboring relay nodes. There may also be some loss in network throughput due to transient looping behavior in the routing of some packets.

It was observed in the studies of ARPA's packet radio network [KAHN 78] that a distributed adaptive algorithm is not suitable for packet radio networks operating in a hostile environment. The proper

Routing algorithm	Throughput efficiency	Availability/ survivability	Handling of mobile users	Large network implications
1. Distributed adaptive	good	poor	poor	large nodal routing tables and processing requirements; long delays
2. Flooding	poor	excellent	excellent	network throughput efficiency becomes worse; long delays
3. Fixed routes				large station routing tables and processing requirements; long delays
(a) single station	good	poor	poor	
(b) multiple stations	good	fair to poor	poor	
4. Fixed routes established via flooding	fair	good	poor	network throughput efficiency becomes worse; long delays

Figure 1. A qualitative comparison of 4 network routing algorithms.

functioning of the routing algorithm depends upon the correct functioning of all nodes. If any one node broadcasts erroneous routing updates, the integrity of the whole network is jeopardized.

As for the special network attributes of interest here, it is unlikely that this algorithm can respond fast enough to handle high-speed mobile users, especially in a large network. Also, when the network is large, with many relay nodes, both processing and memory requirements for managing routing tables within relay nodes are very large. Additionally, source-destination delays will be long due to the presence of many intermediate store-and-forward hops.

Broadcast routing (flooding)

With broadcast routing, each node repeats a newly received packet to all its neighbors. (In a packet radio network in which each node is equipped with an omni-directional antenna, this is typically accomplished with a single transmission.) The following mechanism is needed to avoid the creation of duplicate copies of a packet in a network without bound. Specifically, each packet needs to have a globally unique identifier. Each node needs to have memory space to store identifiers of packets that it has repeated previously. Assuming that each packet can be identified uniquely, every relay in a network will repeat transmission of the packet at most once. Thus, if a network has N relay nodes, each packet will be transmitted exactly N times in the network.

A packet can be uniquely identified by its source node and a sequence number. Since in practice sequence numbers are finite, an additional mechanism is usually necessary to reduce the probability that a node does not recognize a packet that it has already repeated previously. Typically, either a hop count or an age field can be included with the packet. Nodes receiving packets which are "too old" will discard them.

We see that this routing procedure is implemented with a very simple distributed algorithm. Furthermore, the algorithm utilizes no information about network conditions and achieves routing by simply flooding every node in the whole network with a copy of the packet. Here lie both the advantages and disadvantages of the algorithm. The advantages are the excellent survivability of the network and the ease of handling mobile users (high-speed or not). The main disadvantage is the extremely poor throughput efficiency of the network; the transport of one packet requires N transmissions (involving all relay nodes) in the network. With large networks having many store-and-forward hops between a source and a destination, network delays are long.

Fixed routes established by a single station or multiple stations

Consider first the use of a single station that maintains information on the network connectivity and is responsible for assigning routes. Updates on changes in the network connectivity and traffic conditions are forwarded to the station. Fixed routes exist between network nodes and the station, which may be established by, for instance, the labeling procedure described in [GITM 76]. A node desiring communication with another node obtains a route from the station. The complete route may be carried in the header of each packet traversing from the source to the destination node. For large networks with long routes, this is a significant overhead. Alternatively, once a route has been assigned, the source node may send a control packet to set up the routing information in intermediate nodes along the route. This necessitates nodal processing and memory requirements. Also, with a large network, substantial memory and processing capabilities are required at the station.

With this routing algorithm, network survivability is poor since the station is a single vulnerable point. It is also not very responsive to rapid changes in the network connectivity. Handling of high-speed mobile users will also be difficult. When the network is not subject to rapid changes, the network throughput efficiency is good since shortest path routes are used.

The network may be partitioned into regions with a station for each region. For a very large network, multiple stations can efficiently reduce the size of routing tables and the overhead of route computations. With multiple stations, the network's survivability will also be improved somewhat but the handling of high-speed mobile users remains difficult.

Whether a single station or multiple stations are used in a large network, source destination routes have many hops and so the network transit delays for packets are large.

Fixed routes established by flooding

This is a compromise between the previous two routing schemes. When a node wants to communicate with another node, a route is first discovered via flooding. Subsequent packets are then transmitted along the discovered route. The network throughput efficiency of this scheme is considered to be only fair for two reasons: (1) the overhead needed for flooding to establish routes, and (2) the discovered route may not be the best one.

The network availability is considered good. The handling of mobile users is, however, still poor because if a user is moving fast, the discovered route to him becomes useless very quickly.

Assessment of the four routing algorithms

It is clear that a distributed adaptive routing algorithm is not suitable for very large networks since the vulnerability of a single relay affects the integrity of the whole network. We shall not consider it any further.

As for the other three algorithms, none of them can meet all the primary performance requirements with the given special network attributes. Referring to Figure 1, we see that the strengths and weaknesses of routing schemes 2 and 3 are complementary. While routing scheme 4 is kind of a hybrid between 2 and 3, it fails to capture the strengths of them both. Furthermore, in a large network there are many hops in each route. As a result, network transit delays are long in all cases. When fixed routes are used, substantial overheads are necessary for the specification of routes as well as for the management of routing tables.

3. HIERARCHICAL NETWORKS--A PROPOSAL

It is well known that a hierarchical network structure makes it possible to specify routes and maintain routing tables efficiently [KAMO 76]. We also learned from [GERL 80] that both congestion control and end-to-end flow control techniques have hierarchical structures.

Comparing schemes 2 and 3 in Figure 1, we see that scheme 2 is highly survivable because there is no need for any control information (directory or routing table) to be stored anywhere. However, the network throughput efficiency is extremely poor because each packet generates N transmissions in the network, where N is the number of network relays. We would like to retain much of the survivability of scheme 2 without paying too big a price in network throughput efficiency, especially for large networks. A proposal based upon a hierarchical network structure is given below. The network structure and routing strategies are discussed in Section 3.1. Implementation considerations are discussed in Section 3.2. The impact on flow control and other network algorithms is discussed in Section 3.3.

3.1 Network Structure and Routing

Consider the hierarchical network structure illustrated by Figure 2. Suppose there are n classes of nodes. Nodes in each class are further partitioned into groups. Within each group of nodes, say of class i (i=1,...,n), broadcast routing is used. One of the nodes in each group serves as a group leader and is also given a class i+1 status (for i≠n). It is responsible for relaying packets going in and out of the group by communicating with other class i+1 nodes within other groups at the i+1 level.

For example, let there be 1,000 network nodes organized into groups of 10 nodes so that n=3; there are 100 groups of class 1 nodes, 10 groups of class 2 nodes and 1 group of class 3 nodes. To send a packet from a class 1 node in one group to another, the maximum number of transmissions required is 50 as illustrated by Figure 3. (Compare this with 1,000 transmissions when the network has no hierarchical structure.) The price that one pays compared to routing seheme 2 in Figure 1, is that each group leader must have a directory function; he needs to know all the nodes that are under him. This requirement will not be difficult to meet if groups are relatively small.

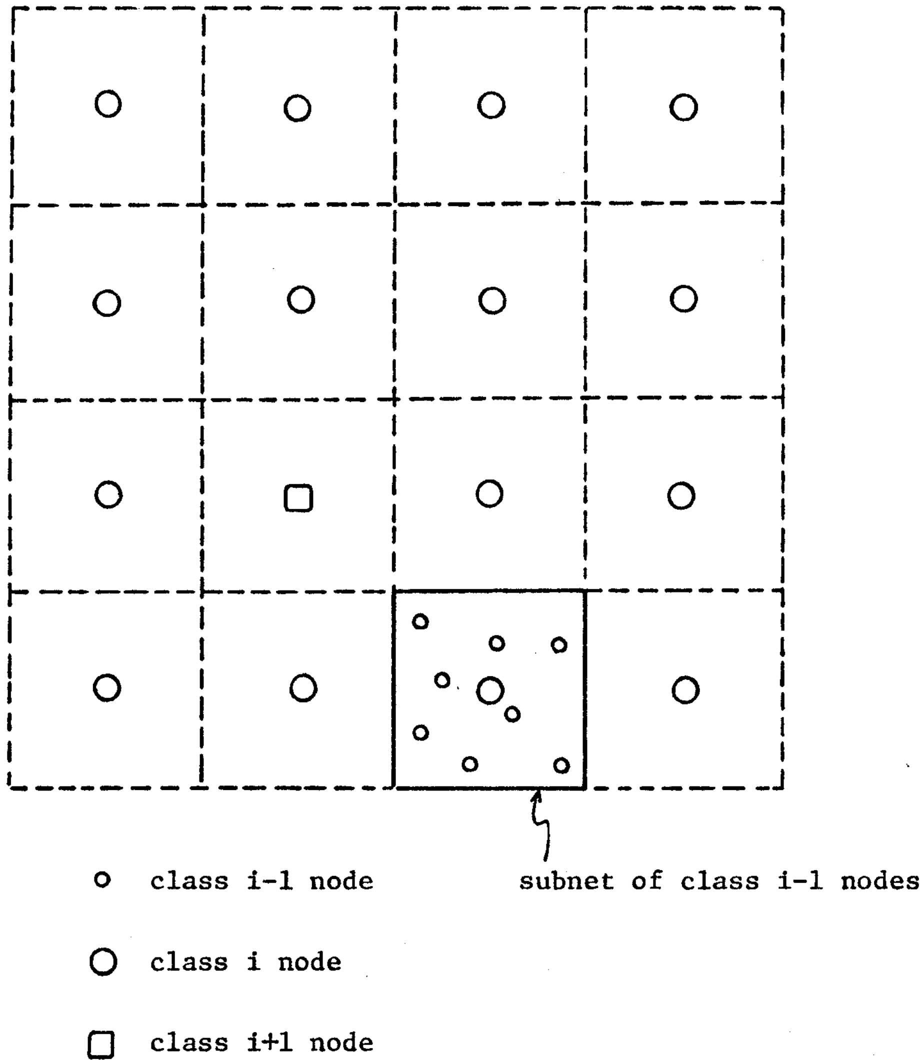

Fig. 2. A hierarchical network structure.

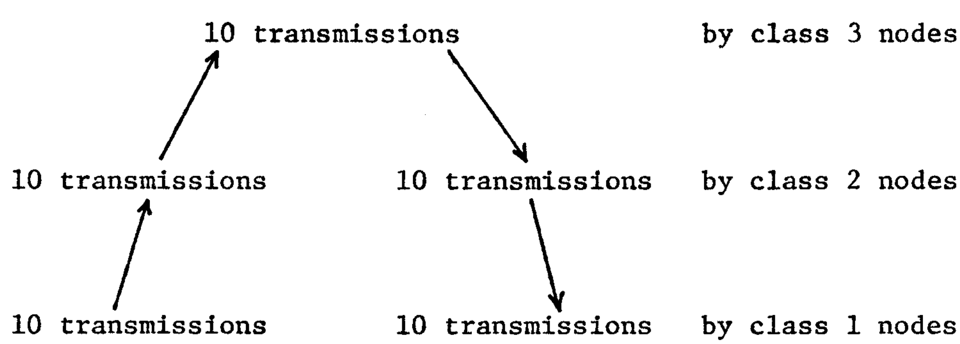

Fig. 3. Number of transmissions needed in a 3-level network.

The differences between this routing strategy and scheme 3 in Figure 1 with multiple stations are: (1) group leaders here have a directory function while stations in scheme 3 have both directory and route assignment functions; (2) the network structure here may have more than two hierarchical levels; (3) we assume that nodes of the same class here have enough transmission power to communicate directly with one another (see below).

We believe that a directory function is easier to maintain and to recover from failures than a route assignment function. If a group leader fails, members of the group can elect another leader. Distributed algorithms to implement such failure recovery functions as well as to handle mobile users that are passing through are conceptually feasible. It is a subject that should be investigated in the future.

What we have proposed above gains throughput efficiency basically by dividing up a network into compartments so that flooding is done one compartment at a time rather than over the whole network.

Distributed algorithms for managing the groups and directories should permit the splitting up and coalescing of groups. In the extreme case, if all groups are coalesced, network routing is achieved via the use of flooding over the entire network (just as in routing scheme 2 of Figure 1).

For network throughput efficiency considerations, it may be desirable to use a flexible combination of fixed routes and broadcast routing. Suppose each group leader maintains both a directory and also attempts to maintain routes to reach nodes under him. Flooding within a group is used only in the event of connectivity changes or for mobile users. For example, each packet may start out by following a fixed route. But if the route length has been exceeded without reaching its destination, broadcast routing is then triggered into effect. Thus, at any one time, a combination of broadcast routing and fixed routes may be used in different parts of the network at different levels.

The above hiererchical network structure and routing strategy will then permit a network to operate with good throughput efficiency using fixed routes when the network connectivity is stable. However, with increasing failures and rapidly changing connectivity, more broadcast routing is used within the network to achieve logical connectivity (thus providing a high degree of availability/survivability). Mobile users are handled by broadcast routing. If they travel at a relatively low speed, then they can be reached by flooding a low level group of nodes. If, on the other hand, they travel at a relatively high speed, then it might be necessary to reach them by flooding a wider geographical area corresponding to a higher level group of nodes.

In summary, reconsidering the performance measures and network attributes listed in Figure 1, the above proposal embodies the strengths of schemes 2 and 3 by attempting to adapt to the changing nature of the network. The adaptivity is made possible with a hierarchical network structure.

3.2 Implementation Considerations

We consider next the network implementation issues of

(1) relay transmission power (range), and
(2) radio channel multiple access.

Relay transmission power

Let us make the usual assumption of a network using a single radio channel (frequency). Each relay uses an omni-directional antenna. One design decision is the transmission power (range) of each relay [KLEI 78, AKAV 78]. If relays have a large transmission power, then they may possibly reach a destination node with a single transmission (except, of course, when line-of-sight obstructions are present). However, each transmission will compete and interfere with a large population of network nodes that lie within the transmission radius. On the other hand, if a small transmission power is used, few network nodes lie within the transmission radius. But now, many hops (transmissions) are needed to deliver a packet from its source node to its destination node. This trade-off was examined by Kleinrock and Silvester [KLEI 78] for a network model with idealized assumptions of slotted ALOHA for multiple access and a form of point-to-point directional routing. They found that the network throughput rate is highest when each relay has a small transmission power with only 5-6 relays within its range (which is also the minimum number required for a connected network). In this analysis, the important consideration of source-destination delay was not considered. If a channel access protocol such as CSMA is used which significantly reduces the effect of multiple access interference (relative to slotted ALOHA), then the source-destination delay is roughly proportional to the number of transmission hops from source to destination. Therefore, when delay is an important performance measure, a large transmission power should be used. A short delay also implies small buffering requirements needed for network relays. (The trade-off between throughput, delay and transmission power has been addressed in [AKAV 78].)

In short, the trade-off is: a large transmission power is preferred to minimize delay and a small transmission power is preferred to maximize throughput. For a hierarchical network such as proposed in Section 3.1 above, and illustrated in Figure 2, we advocate the following strategy for designing relay transmission power. Nodes of the same class and belonging to the same group should have enough transmission power to form a connected network (each such group will be referred to as a subnet). Thus, referring to Figure 2, a transmission that emanates from a class i+1 node will have enough power to reach its adjacent class i+1 nodes, without having to be relayed by class i nodes that lie in between. This strategy is necessary to take advantage of the hierarchical network structure to have a small number of hops between a source node and a destination node and consequently a small network delay.

The following variation may also be used for low level nodes that belong to the same subnet which spans only a small geographical area. The transmission power of each node may be sufficient for every node to reach every other node within the subnet in one hop (except for line-of-sight obstructions) so that store-and-forward routing, either fixed routes or flooding, is not necessary.

Depending upon the geographical distances involved in the network, it may occasionally happen that some high level nodes are too far apart to transmit to each other directly. Some possible solutions are: (1) use of directional antennas for transmission, (2) these nodes may communicate via multi-hop paths over nodes at the next lower level in the network hierarchy (this is somewhat like ARPA's packet radio network [KAHN 78]), and (3) use of an air-borne relay.

Since high level nodes use a large transmission power, it may be desirable to provide separate radio channels (frequencies) for the different classes of nodes. Alternatively, a priority CSMA protocol, to be described below, may be used.

Channel access

It is beyond the scope of this report to consider the use of a spread spectrum technique for security and anti-jam reasons [KAHN 78]. A spread spectrum technique, if employed, is assumed to have been implemented as part of the radio channel under consideration. Our concern is confined to the sharing of the radio channel for throughput-delay performance considerations. In this respect, assuming relatively small geographical distances and thus small channel propagation delays, CSMA protocols provide the best performance among multiple access protocols [KLEI 75a, LAM 80a].

One way to implement the hierarchical network structure proposed above is to provide separate radio channels for the different classes of nodes. CSMA is used for accessing each channel.

If only a single radio channel is available for all classes of nodes, a CSMA protocol with preemptive priority may be used [LAM 80b]. This protocol is illustrated by Figure 4. Let class n nodes have the highest priority, while class 1 nodes the lowest priority. Suppose nodes in classes from i+1 to n have nothing to send. Following a class i packet transmission in the radio channel, class i nodes will have the first priority to access the channel among nodes in classes 1, 2,...,i. Class i-1 nodes will defer for τ seconds and will then attempt channel access if the channel is idle at that time.

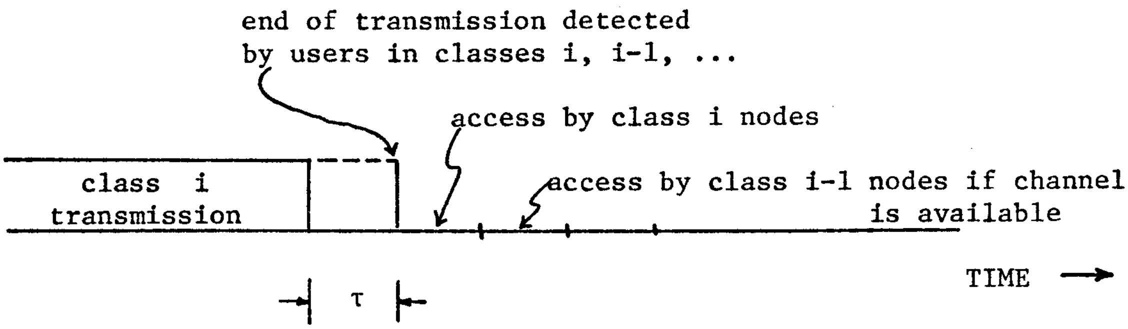

Fig. 4. An illustration of a CSMA protocol with priority access.

In general, class i nodes will do carrier sensing and defer to an ongoing transmission only for other nodes in the i-th or higher classes. Since it may not be able to hear the transmissions of nodes in lower classes (due to their smaller transmission power), its transmission may interfere with transmissions by nodes in lower classes. When that happens, the lower class transmissions will be corrupted by the collision but there is a good chance that the class i transmission will be received correctly because of its larger power. (This is the FM capture problem discussed in [KAHN 78].) In any case, after the collision, all other nodes in classes 1 to i will defer to this class i transmission. It can then access the channel unless, of course, it is preempted in turn by a transmission from a higher class.

We note that the above CSMA preemptive priority protocol will not achieve as high a network throughput rate as a pure CSMA protocol for two reasons: (1) long transmission range used by the high level nodes [KLEI 78], and (2) carrier sensing is performed only for other nodes in the same or a higher class. The trade-off is to achieve a smaller network delay by reducing the number of hops along routes.

3.3 Impact on Flow Control and Other Network Algorithms

Apart from improving network delay, a hierarchical structure is desirable for a large network in many other ways. These are discussed below.

Routing

If fixed routing is used, the size of routing tables is significantly reduced [KAMO 76]. If source routing is used with the entire route specified in the header of each packet, the overhead for representing a route in a packet header is substantially less since routes have fewer hops.

Store-and-forward buffer deadlock avoidance

To avoid store-and-forward deadlocks using the buffer allocation strategy discussed in [GUEN 75], a necessary condition is that each relay node must have more buffers than the largest number of hops that an arriving packet has traversed. With a large network, say 1,000 relay nodes, the number of buffers required at each relay node for deadlock avoidance is extremely large. A hierarchical network structure makes it possible to avoid store-and-forward deadlocks with a reasonable number of buffers at each relay node ($\geq$ number of hops traversed by a packet within relay nodes in the same subnet).

Congestion control

The various network congestion control techniques discussed in [GERL 80] will all benefit from a hierarchical network structure. Consider the input buffer limit scheme [LAM 79, LAM 80c]. For it to be effective, we must have

$$\frac{N_I}{N_T} < \frac{1}{H}$$

where N_I is the number of input buffers in a relay node, and N_T is the total number of buffers, and H is the mean number of hops traversed by a packet in the network. The above rule assumes the use of fixed routes. If flooding is used instead, then the above rule should be modified to

$$\frac{N_I}{N_T} < \frac{1}{N}$$

where N is the number of relay nodes in the network. Thus we see that with a large network, the provision of even a single input buffer would give rise to a very large buffer requirement at each relay node. To reduce the number of buffers required, one can further restrict the use of an input buffer to a part-time basis. For example, a packet awaiting an end-to-end reply from its destination node occupies the input buffer and blocks further input. This may be too much of a restriction on the network's input rate.

With a hierarchical structure, the input buffer limit strategy may be implemented for each subnet. The overall network congestion control problem is then decomposed into a series of congestion control problems for smaller communities of nodes (subnets). The buffer requirements will be substantially less because of the fewer number of nodes in each subnet. The trade-off is that now the network is lossy, i.e. a packet may be refused admittance to a subnet at a different level and will have to be discarded by its current node. The network error control protocols will have to be designed with this possibility in mind (to be discussed below).

Both the isarithmic technique [DAVI 72] and choke packet scheme [MAJI 79] for network congestion control will also not perform well in a large network. With the isarithmic technique, the main difficulty is the adaptive distribution of permits to places that need them, which will be even harder to do in a large network. The responsiveness of of the choke packet scheme to local congestions will be affected in a large network; the propagation of choke packets to the entry points of virtual channels will require a longer delay. Both schemes can be adapted to a hierarchical network in a manner similar to the above described for the input buffer limit strategy. In general, a hierarchical network structure allows the problem of managing a large network to be decomposed into smaller, more manageable problems.

End-to-end flow control

To perform source-sink speed synchronization over a large number of intermediate nodes is difficult due to the large delay involved. Also, a very large window will be necessary to maintain the throughput rate of a virtual channel at an acceptable level. Large virtual channel windows increase the difficulty of network congestion control and the management of buffers to implement functions for the sequencing of packets and the reassembly of packets into messages. A solution is to achieve the overall end-to-end flow control with a series of concatenated end-to-end controls as discussed in [GERL 80]. A hierarchical network such as described above provides a natural framework for doing so.

Error control

In packet radio networks, it is desirable that the store-and-forward protocol of relay nodes permits the discarding of packets that cannot be forwarded after a number of tries. This is necessary to implement adaptive control for the contention access of a radio broadcast channel to prevent instability behavior [LAM 75, LAM 80a].

We said earlier that to provide deadlock-free buffering for a large network requires an extremely large number of buffers at each relay node. One way to provide deadlock freedom without the use of a large number of buffers at each relay node is to allow the network to be occasionally lossy.

Another solution, described earlier, is to use a hierarchical network structure to reduce the number of buffers required at relay nodes. If a hierarchical network structure is used and congestion control is performed for individual subnets at different levels, then the network is again lossy; occasionally, a packet may be refused entry into a subnet at a different level and its current node has no buffer to keep it any longer and must discard it.

If routing is done by flooding, the discarding of a copy of a packet will not likely cause any problem. Whichever routing strategy is used, however, error control needs to be provided between the source and the destination so that if a packet is lost by the network it can be retransmitted from its source. For a very large network with a hierarchical structure, it may be desirable to introduce error control between intermediate recovery points. This provision avoids having to recover a lost packet all the way back from its source; source-destination error control, however, should still be used to provide a highly reliable end-to-end communication.

Fairness and priority allocation of network capacity

A mechanism to control the allocation of network "capacity" to the users of different virtual channels is the window size (number of packets allowed in transit within the network) of a virtual channel. It was found that if virtual channels are given window sizes proportional to their path lengths, then the variance in their throughput rates is quite small. This may be considered as a fair allocation policy [LAM 80c].

In large networks, large window sizes will need to be used to maintain an acceptable level of throughput rate for virtual channels. Due to the increased variability of network delays, the end-to-end windows will become less effective in controlling the allocation of network capacity. Additional controls such as step-wise control (concatenated end-to-end windows) discussed in [GERL 80] may be needed.

It is important to be able to provide special packets or virtual channels with priority classes of delay service. There are several places in the network architecture where priority scheduling may be used to improve the network transit delay of important packets. A priority queueing discipline may be implemented at each relay node. A priority scheme may be used in the acceptance or rejection of packets in conjunction with congestion control and deadlock avoidance protocols. Note that if a priority CSMA protocol is used, then high level nodes will have priority preemptive access to the radio channel; the network hierarchy should therefore reflect the rank (importance) hierarchy of the network's user community.

With a network that may lose packets, retransmission from a source or even a recovery point may result in an unacceptable delay for some very important packets. In this case, the technique of dispersity routing may be used for these messages [MAXE 75]. An error correction code, say a (7, 4) Hamming code, is used to transform each message into seven packets. Any 4 of the 7 packets received at the destination will be sufficient for the original message to be reconstructed. This technique thus improves both reliability and network transit delay at the expense of some network throughput efficiency. It should therefore only be used for very high priority messages.

4. CONCLUSIONS

We considered the design of large packet radio networks with the key attributes: (1) large population of users, (2) mobile users, and (3) changing network connectivity. Four routing algorithms were examined. We found that none of these algorithms are satisfactory with respect to the primary performance measures of throughput efficiency, network delay and availability/survivability given the above network attributes. We then proposed a conceptual architecture based upon a hierarchical network structure and routing strategy which appears to be able to combine the good features of the above routing algorithms to meet all of the primary performance requirements. Additional research will be needed to develop distributed algorithms for network topology maintenance and routing.

The implementation of such a network architecture requires several classes of packet radios with different transmission power ranges. A CSMA channel access protocol with preemptive priority was proposed. The impact of the network architecture on the design of network algorithms for flow control and error control were also addressed.

REFERENCES

AKAV 78 Akavia, G. Y., "Hierarchical Organization of Distributed Packet-switching Communication Systems," Ph.D. Dissertation, 1978, Department of Computer Sciences, UCLA, Los Angeles.

DAVI 72 Davis, D. W., "The Control of Congestion in Packet Switching Networks," IEEE Trans. on Commun., Vol COM-20, June 1972.

FRAN 75 Frank, H., I. Gitman and R. Van Slyke, "Packet Radio System Network Considerations," Proc. Nat. Comput. Conf., May 1975.

GERL 80 Gerla, M. and L. Kleinrock, "Flow Control: A Comparative Survey," IEEE Trans. on Commun., April 1980.

GITM 76 Gitman, I. et al., "Routing in Packet-switched Broadcast Radio Networks," IEEE Trans. on Commun., August 1976.

GUEN 75 Guenther, K. D., "Prevention of Buffer Deadlocks in Packet-switching Networks," IFIP-IIASA Workshop on Data Communications, Luxenburg, Austria, 1975.

KAHN 77 Kahn, R. E., "The Organization of Computer Resources into a Packet Radio Network," IEEE Trans. Commun., Vol. COM-25, No. 1. January 1977.

KAHN 78 Kahn, R. E. et al., "Advances in Packet Radio Technology," Proc. IEEE, November 1978.

KAMO 76 Kamoun, F., "Design Considerations for Large Computer Communications Networks," Ph.D. dissertation, Univ. of California, Los Angeles, Eng. Rep. 7642, April 1976.

KLEI 75a Kleinrock, L. and F. Tobagi, "Packet Switching in Radio Channels: Part I--Carrier Sense Multiple Access Modes and Their Throughput-delay Characteristics," IEEE Trans. Commun., December 1975.

KLEI 78 Kleinrock, L. and J. Silvester, "Optimum Transmission Radii for Packet Radio Networks or Why Six is a Magic Number," Conference Record of the IEEE National Telecommunication Conference, Birmingham, Alabama, December 3-6 1978.

LAM 75 Lam, S. S. and L. Kleinrock, "Dynamic Control Schemes for a Packet Switched Multi-access Broadcast Channel," AFIPS Conf. Proc., Vol. 44, 1975.

LAM 79 Lam, S. S. and M. Reiser, "Congestion Control of Store-and-forward Networks by Input Buffer Limits -- An Analysis," IEEE Trans. on Commun., January 1979.

LAM 80a Lam, S. S., "A Carrier Sense Multiple Access Protocol for Local Networks," Computer Networks, 1980.

LAM 80b Lam, S. S., "A Packet Network Architecture for Local Interconnection," Conf. Rec. ICC, Seattle, June 1980.

LAM 80c Lam, S. S. and Y. L. Lien, "An Experimental Study of the Congestion Control of Packet Communication Networks," Proc. Fifth ICCC, Atlanta, October 1980.

MAJI 79 Majithia, J. C. et al., "Experiments in Congestion Control Techniques," in Proc. Int. Symp. Flow Control Comput. Networks, Versailles, France, February 1979.

MAXE 75 Maxemchuk, N. F., "Dispersity Routing," Conf. Rec. ICC, San Francisco, June 1975.

MCQU 80 McQuillan, J. et al., "The New Routing Algorithm for the ARPANET," IEEE Trans. on Commun., May 1980.